

Updated 05/02/2025

Sign up

ElasticSearch Ops training: administration and infrastructure

2 days (14 hours)

Presentation

ElasticSearch in infrared and production mode, use all its power with our advanced training.

Have you attended our [Elasticsearch](#) training course, and would like to know more about ElasticSearch administration? Would you like to learn about best practices in the day-to-day management of ES clusters? This training course will enable you to become an expert in Elasticsearch administration and infrastructure.

In this Elasticsearch training course, we'll look at how to install and configure Elasticsearch in a high-density, production-critical operating environment.

Next, we'll take a look at cluster management and monitoring. Finally, we'll learn about best practices in case of failure, and how to protect your data with X-Pack Security.

The will use the latest stable version of the project ([Elasticsearch 8.17](#)).

Objectives

- Managing a cluster, from installation to operational monitoring on ElasticSearch
- Best practices for ElasticSearch outages
- Protecting your data on ElasticSearch

Target audience

- System Administrator
- Developer
- Architect
- Project Manager
- Big Data Engineer

Prerequisites

Knowledge of operating system administration, and basic knowledge of Elasticsearch.

ElasticSearch 8: Administration and Infrastructure training program

Background on Elasticsearch

Installing, configuring and implementing Elasticsearch

- Apache Lucene
- Scoring and query relevance
- Data import
- What's new in versions 6, 7 & 8

Search and index configuration

- Data indexing, advanced searches & percolation
- Extraction of unstructured data
- Large query and indexing volumes

Clustering

- Knot, Sharding
- Distributed nature of ElasticSearch: Storage & Computation
- Fault tolerance
- Master node concept
- Primary shard and replication
- Internal request routing

Cluster management: configuration & management

- Node configuration & Template management
- The elasticsearch.yml file
- Query tools & /_cluster/health
- Creating an index
- Shard definition
- Node allocation & Configuring new nodes

Case of failure

- Hardware fault tolerance: adjust & optimize storage distribution
- Operation in the event of node loss
- Electing a new master node
- Declaration of new primary shards
- Backup and restoration

Monitoring: Cluster monitoring (formerly Marvel)

- Collecting data
- Operation: log management
- API: administration & supervision
- Cluster monitoring
- Load distribution
- Node allocation & Index aliases
- Transient" and "Persistent" properties

X-Pack Security: protecting your data (formerly Shield)

- Manage users and roles
- Prevent intrusions, data tampering and sniffing
- Track all activities

ADDITIONAL MODULE IN ENGLISH ON REQUEST (+2 DAYS)

- Training language : English
- Course level : Beginner to intermediate

This training course enables you to master the basic concepts of Elasticsearch and to efficiently design, size and manage a cluster. It covers best practices in terms of operability, performance optimization and common problems encountered. Theory: 60% Practical: 40% Audience :

- Data Engineers
- Architects
- System Administrators
- DevOps

Prerequisites :

- System Knowledge required
- Knowledge of REST/HTTP, Json, Yaml are appreciated
- Linux or Mac OS

Getting Started

- Elasticsearch Overview
- Keys Features
- Basic Concepts
- Install Elasticsearch
- CRUD Operations
- First steps on Search API

Configuring Elasticsearch

- Elasticsearch Directories
- Index settings
- Node and Cluster settings
- JVM settings

Mapping and Analysis

- Introduction
- Data Types
- Main parameters
- API mapping
- Analysis and Inverted Index
- Custom Analyzer
- Multi-Fields

Index Management

- Index Configuration
- Alias
- Templates

Nodes and Cluster Management

- Node Types
- Cluster State
- Master Node
- Data Node
- Coordinating Only Node
- Ingest Node

Shard Management

- Health Cluster
- Shard Allocation
- Shard state

Inside Node

- Segment
- Merging Segment
- Elasticsearch indexing
- Elasticsearch Storage

Caches in Elasticsearch

- Capacity Planning
- Designing for Scale
- Use Cases

Monitoring

- API Stats
- Task API
- Cat API
- Monitoring Component
- Watcher

Architectures

- Hot / Warm Architecture
- Multi-Cluster Architecture
- Elastic Stack and Kafka Integration

Companies concerned

This course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced computer technology, or to acquire specific business knowledge or modern methods.

Positioning on entry to training

Positioning at the start of training complies with Qualiopi quality criteria. As soon as registration is finalized, the learner receives a self-assessment questionnaire which enables us to assess his or her estimated level of proficiency in different types of technology, as well as his or her expectations and personal objectives for the training to come, within the limits imposed by the selected format. This questionnaire also enables us to anticipate any connection or security difficulties within the company (intra-company or virtual classroom) which could be problematic for the follow-up and smooth running of the training session.

Teaching methods

Practical course: 60% Practical, 40% Theory. Training material distributed in digital format to all participants.

Organization

The course alternates theoretical input from the trainer, supported by examples, brainstorming sessions and group work.

Validation

At the end of the session, a multiple-choice questionnaire verifies the correct acquisition of skills.

Sanction

A certificate will be issued to each trainee who completes the course.